

Probabilistic coherence and proper scoring rules*

Joel Predd Rand Corporation	Robert Seiringer Princeton University	Elliott H. Lieb Princeton University
Daniel Osherson Princeton University	Vincent Poor Princeton University	Sanjeev Kulkarni Princeton University

October 16, 2007

Abstract

We provide self-contained proof of a theorem relating probabilistic coherence of forecasts to their non-domination by rival forecasts with respect to any proper scoring rule. The theorem appears to be new but is closely related to results achieved by other investigators.

1 Introduction

Scoring rules measure the quality of a probability-estimate for a given event, with lower scores signifying probabilities that are closer to the event's status (1 if it occurs, 0 otherwise). The sum of the scores for estimates \mathbf{p} of a vector \mathcal{E} of events is called the “penalty” for \mathbf{p} . Consider two potential defects in \mathbf{p} .

- There may be rival estimates \mathbf{q} for \mathcal{E} whose penalty is guaranteed to be lower than the one for \mathbf{p} , regardless of which events come to pass.
- The events in \mathcal{E} may be related by inclusion or partition, and \mathbf{p} might violate constraints imposed by the probability calculus (for example, that the estimate for an event not exceed the estimate for any event that includes it).

*© 2007 by the authors. This paper may be reproduced, in its entirety, for non-commercial purposes. Research supported by NSF grants PHY-0652854 to Lieb, and PHY-0652356 to Seiringer. R.S. also acknowledges partial support from an A. P. Sloan Fellowship.

Building on the work of earlier investigators (see below), we show that for a broad class of scoring rules known as “proper” the two defects are equivalent. An exact statement appears as Theorem 1. To reach it, we first explain key concepts intuitively (the next section) then formally (Section 3). Proof of the theorem proceeds via three propositions of independent interest (Section 4). We conclude with generalizations of our results and an open question.

2 Intuitive account of concepts

Imagine that you attribute probabilities .6 and .9 to events E and F, respectively, where $E \subseteq F$. It subsequently turns out that F comes to pass but not E. How shall we assess the perspicacity of your two estimates, which may jointly be called a **probabilistic forecast**? According to one method (due to Brier, 1950) truth and falsity are coded by 1 and 0, and your estimate of the chance of E is assigned a score of $(0 - .6)^2$ since E did not come true (so your estimate should ideally have been zero). Your estimate for F is likewise assigned $(1 - .9)^2$ since it should have been one. The sum of these numbers serves as overall **penalty**.

Let us calculate your expected penalty for E (prior to discovering the facts). With .6 probability you expected a score of $(1 - .6)^2$, and with the remaining probability you expected a score of $(0 - .6)^2$, hence your overall expectation was $.6(1 - .6)^2 + .4(0 - .6)^2 = .24$. Now suppose that you attempted to improve (lower) this expectation by insincerely announcing .65 as the chance of E, even though your real estimate is .6. Then your expected penalty would be $.6(1 - .65)^2 + .4(0 - .65)^2 = .2425$, worse than before. Differential calculus reveals the general fact:

Suppose your probability for an event E is p, that your announced probability is x, and that your penalty is assessed according to the rule: $(1 - x)^2$ if E comes out true; $(0 - x)^2$ otherwise. Then your expected penalty is uniquely minimized by choosing $x = p$.

Our scoring rule thus encourages sincerity since your interest lies in announcing probabilities that conform to your beliefs. Rules like this are called **proper**. (We add a continuity condition in our formal treatment, below.) For an example of an improper rule, substitute absolute deviation for squared deviation in the original scheme. According to the new rule, your expected penalty for E is $.6|1 - .6| + .4|0 - .6| = .48$ whereas it drops to $.6|1 - .65| + .4|0 - .65| = .47$ if you fib as before.

Consider next the rival forecast of .95 for E and .55 for F. Because $E \subseteq F$, this forecast is inconsistent with the probability calculus (or **incoherent**). Table 1 shows that the original forecast **dominates** the rival inasmuch as its penalty is lower however the facts play out. This association of incoherence and domination is not an accident. No matter what proper scoring rule is in force, any incoherent forecast can be replaced by a coherent one whose penalty is lower in every possible circumstance; there is no such replacement for a coherent forecast. This fact is formulated as Theorem 1 in the next section. It can be seen as partial vindication of probability as an expression of chance.¹

Logical possibilities	Forecast	
	original	rival
E = T F = T	.17	.205
E = F F = T	.37	1.105
E = F F = F	1.17	1.205

Table 1: Penalties for two forecasts in alternative possible realities

These ideas have been discussed before, first by de Finetti (1974) who began the investigation of dominated forecasts and probabilistic consistency (called **coherence**). His work relied on the **quadratic scoring rule**, introduced above.² Lindley (1982) generalized de Finetti’s theorem to a broad class of scoring rules. Specifically, he proved that for every sufficiently regular generalization s of the quadratic score, there is a transformation $T : \mathfrak{R} \rightarrow \mathfrak{R}$ such that a forecast \mathbf{f} is not dominated by any other forecast with respect to s if and only if the transformation of \mathbf{f} by T is probabilistically coherent. The reliance on the transformation T , however, clouds the interpretation of Lindley’s theorem.

Fresh insight into proper scoring rules comes from relating them to a generalization of metric distance known as **Bregman divergence** (Bregman, 1967). This relationship was studied by Savage (1971), albeit implicitly, and more recently by Banerjee et al. (2005) and Gneiting and Raftery (2007). So far as we know, their results have yet to be connected to the issue of dominance.

To pull together the threads of earlier discussions, the present work offers a self-contained account of the relations among (i) coherent forecasts, (ii) Bregman divergences, and (iii) domination with respect to proper scoring rules. Only elementary analysis is presupposed. We begin by formalizing the concepts introduced above.³

¹The other classic vindication involves sure-loss contracts; see Skyrms (2000).

²For analysis of de Finetti’s work, see Joyce, 1998. Note that some authors use the term **inadmissible** to qualify dominated forecasts.

³For application of scoring rules to the assessment of opinion, see Gneiting and Raftery (2007) along with Bernardo and Smith (1994, §2.7.2) and references cited there.

3 Framework and Main Result

Let Ω be a nonempty **sample space**. Subsets of Ω are called **events**. Let \mathcal{E} be a vector (E_1, \dots, E_n) of $n \geq 1$ events over Ω . We assume that Ω and \mathcal{E} have been chosen and are now fixed for the remainder of the discussion. We require \mathcal{E} to have finite dimension n but otherwise our results hold for any choice of sample space and events. In particular, Ω can be infinite. We rely on the usual notation $[0, 1]$, $(0, 1)$, $\{0, 1\}$ to denote, respectively, the closed interval $\{x : 0 \leq x \leq 1\}$, the open interval $\{x : 0 < x < 1\}$ and the two-point set containing 0, 1.

Definition 1. Any element of $[0, 1]^n$ is called a **(probability) forecast (for \mathcal{E})**. A forecast \mathbf{f} is **coherent** just in case there is a probability measure μ over Ω such that for all $i \leq n$, $f_i = \mu(E_i)$.

A forecast is thus a list of n numbers drawn from the unit interval. They are interpreted as claims about the chances of the corresponding events in \mathcal{E} . The first event in \mathcal{E} is assigned the probability given by the first number (f_1) in \mathbf{f} , and so forth. A forecast is coherent if it is consistent with some probability measure over Ω .

This brings us to scoring rules. In what follows, the numbers 0 and 1 are used to represent falsity and truth, respectively.

Definition 2. A function $s : \{0, 1\} \times [0, 1] \rightarrow [0, \infty]$ is said to be a **proper scoring rule** in case

- (a) $ps(1, x) + (1 - p)s(0, x)$ is uniquely minimized at $x = p$ for all $p \in [0, 1]$.
- (b) s is continuous, meaning that for $i \in \{0, 1\}$, $\lim_{n \rightarrow \infty} s(i, x_n) = s(i, x)$ for any sequence $x_n \in [0, 1]$ converging to x .

For condition 2(a), think of p as the probability you have in mind, and x as the one you announce. Then $ps(1, x) + (1 - p)s(0, x)$ is your expected score. Fixing p (your genuine belief), the latter expression is a function of the announcement x . Proper scoring rules encourage candor by minimizing the expected score exactly when you announce p .⁴

The continuity condition is consistent with s assuming the value $+\infty$. This can only occur for the arguments $(0, 1)$ or $(1, 0)$, representing categorically mistaken judgment. For if $s(0, p) = \infty$ for some $p \neq 1$, then $ps(1, x) + (1 - p)s(0, x)$ can not have a unique minimum at $x = p$; similarly, $s(1, p) < +\infty$ for $p \neq 0$. A typical example of an unbounded proper scoring rule is $s(i, x) = -\ln|i - x|$ (Good, 1952). A comparison of alternative rules is offered in Selten (1998).

⁴Some authors call such rules *strictly proper*.

For an event E , we let C_E be the characteristic function of E ; that is, for all $\omega \in \Omega$, $C_E(\omega) = 1$ if $\omega \in E$ and 0 otherwise. Intuitively, $C_E(\omega)$ reports whether E is true or false if Nature chooses ω .

Definition 3. Given proper scoring rule s , the **penalty** P_s based on s for forecast \mathbf{f} and $\omega \in \Omega$ is given by:

$$P_s(\omega, \mathbf{f}) = \sum_{i \leq n} s(C_{E_i}(\omega), f_i). \quad (1)$$

Thus, P_s sums the scores (conceived as penalties) for all the events under consideration. Henceforth, the proper scoring rule s is regarded as given and fixed. The theorem below holds for any choice we make.

Definition 4. Let a forecast \mathbf{f} be given.

- (a) \mathbf{f} is **weakly dominated** by a forecast \mathbf{g} in case $P_s(\omega, \mathbf{g}) \leq P_s(\omega, \mathbf{f})$ for all $\omega \in \Omega$.
- (b) \mathbf{f} is **strongly dominated** by a forecast \mathbf{g} in case $P_s(\omega, \mathbf{g}) < P_s(\omega, \mathbf{f})$ for all $\omega \in \Omega$.

Strong domination by a rival, coherent forecast \mathbf{g} is the price to be paid for an incoherent forecast \mathbf{f} . Indeed, we shall prove:

Theorem 1. Let a forecast \mathbf{f} be given.

- (a) If \mathbf{f} is coherent then it is not weakly dominated by any forecast $\mathbf{g} \neq \mathbf{f}$.
- (b) If \mathbf{f} is incoherent then it is strongly dominated by some *coherent* forecast \mathbf{g} .

Thus, if \mathbf{f} and \mathbf{g} are coherent and $\mathbf{f} \neq \mathbf{g}$ then neither weakly dominates the other. The theorem follows from three propositions of independent interest, stated in the next section. We close the present section with a corollary.

Corollary 1. A forecast \mathbf{f} is weakly dominated by a forecast $\mathbf{g} \neq \mathbf{f}$ if and only if \mathbf{f} is strongly dominated by a coherent forecast.

Proof of Corollary 1. The right-to-left direction is immediate from Definition 4. For the left-to-right direction, suppose forecast \mathbf{f} is weakly dominated by some $\mathbf{g} \neq \mathbf{f}$. Then by Theorem 1(a), \mathbf{f} is not coherent. So by Theorem 1(b), \mathbf{f} is strongly dominated by some coherent forecast. \square

4 Three Propositions

The first proposition is a characterization of coherence. It is due to de Finetti (1974).

Definition 5. Let $V = \{(C_{E_1}(\omega), \dots, C_{E_n}(\omega)) : \omega \in \Omega\} \subseteq \{0, 1\}^n$. Let the cardinality of V be k . Let $\text{conv}(V)$ be the **convex hull** of V , i.e., $\text{conv}(V)$ consists of all vectors of form $\alpha_1 \mathbf{v}_1 + \dots + \alpha_k \mathbf{v}_k$, where $\mathbf{v}_i \in V$, $\alpha_i \geq 0$, and $\sum_{i=1}^k \alpha_i = 1$.

The E_i may be related in various ways, so $k < 2^n$ is possible (indeed, this is the case of interest).

Proposition 1. A forecast \mathbf{f} is coherent if and only if $\mathbf{f} \in \text{conv}(V)$.

The next proposition characterizes scoring rules in terms of convex functions. Recall that a convex function φ on a convex subset of \mathfrak{R}^n satisfies $\varphi(\alpha \mathbf{x} + (1 - \alpha)\mathbf{y}) \leq \alpha \varphi(\mathbf{x}) + (1 - \alpha)\varphi(\mathbf{y})$ for all $0 < \alpha < 1$ and all \mathbf{x}, \mathbf{y} in the subset. *Strict* convexity means that the inequality is strict unless $\mathbf{x} = \mathbf{y}$. Variants of the following fact are proved in Savage (1971), Banerjee et al. (2005), and Gneiting and Raftery (2007).

Proposition 2. Let s be a proper scoring rule. Then the function $\varphi : [0, 1] \rightarrow \mathfrak{R}$ defined by $\varphi(x) = -xs(1, x) - (1 - x)s(0, x)$ is a bounded, continuous and strictly convex function, differentiable for $x \in (0, 1)$. Moreover,

$$s(\mathbf{i}, x) = -\varphi(x) - \varphi'(x)(\mathbf{i} - x) \quad \forall x \in (0, 1). \quad (2)$$

Conversely, if a function s satisfies (2), with φ bounded, strictly convex and differentiable on $(0, 1)$, and s is continuous on $[0, 1]$, then s is a proper scoring rule.

We note that the right side of (2), which is only defined for $x \in (0, 1)$, can be continuously extended to $x = 0, 1$. This is the content of the Lemma 1 in the next section. If the extended s satisfies (2) then:

$$s(0, 0) = -\varphi(0) \quad \text{and} \quad s(1, 1) = -\varphi(1). \quad (3)$$

Finally, our third proposition concerns a well known property of Bregman divergences (see, e.g., Censor and Zenios, 1997). When we apply the proposition to the proof of Theorem 1, C will be the unit cube in \mathfrak{R}^n .

Definition 6. Let C be a convex subset of \mathfrak{R}^n with non-empty interior. Let $\Phi : C \rightarrow \mathfrak{R}$ be a strictly convex function, differentiable in the interior of C , whose gradient $\nabla \Phi$ extends to a bounded, continuous function on C . For $\mathbf{x}, \mathbf{y} \in C$, the **Bregman divergence** $d_\Phi : C \times C \rightarrow \mathfrak{R}$ corresponding to Φ is given by

$$d_\Phi(\mathbf{y}, \mathbf{x}) = \Phi(\mathbf{y}) - \Phi(\mathbf{x}) - \nabla \Phi(\mathbf{x}) \cdot (\mathbf{y} - \mathbf{x}).$$

Because of the strict convexity of Φ , $d_\Phi(\mathbf{y}, \mathbf{x}) \geq 0$ with equality if and only if $\mathbf{y} = \mathbf{x}$.

Proposition 3. Let $d_\Phi : \mathcal{C} \times \mathcal{C} \rightarrow \mathfrak{R}$ be a Bregman divergence, and let $Z \subseteq \mathcal{C}$ be a closed convex subset of \mathfrak{R}^n . For $\mathbf{x} \in \mathcal{C} \setminus Z$, there exists a unique $\boldsymbol{\pi}_\mathbf{x} \in Z$, called the **projection of \mathbf{x} onto Z** , such that

$$d_\Phi(\boldsymbol{\pi}_\mathbf{x}, \mathbf{x}) \leq d_\Phi(\mathbf{y}, \mathbf{x}) \quad \forall \mathbf{y} \in Z.$$

Moreover,

$$d_\Phi(\mathbf{y}, \boldsymbol{\pi}_\mathbf{x}) \leq d_\Phi(\mathbf{y}, \mathbf{x}) - d_\Phi(\boldsymbol{\pi}_\mathbf{x}, \mathbf{x}) \quad \forall \mathbf{y} \in Z, \mathbf{x} \in \mathcal{C} \setminus Z. \quad (4)$$

Its worth observing that Proposition 3 also holds if $\mathbf{x} \in Z$, in which case $\boldsymbol{\pi}_\mathbf{x} = \mathbf{x}$ and (4) is trivially satisfied.

5 Proof of Theorem 1

The main idea of the proof is more apparent when s is bounded. So we consider this case on its own before allowing s to reach $+\infty$.

Bounded Case.

Suppose s is bounded. In this case, the derivative of the corresponding φ from Eq. (2) in Proposition 2 is continuous and bounded all the way up to the boundary of $[0, 1]$.

Let $\mathbf{f} \in [0, 1]^n$ be a forecast and, for $\omega \in \Omega$, let $\mathbf{v}_\omega \in V$ be the vector with components $C_{E_i}(\omega)$. Let $\Phi(\mathbf{x}) = \sum_{i=1}^n \varphi(x_i)$. Then

$$\begin{aligned} P_s(\omega, \mathbf{f}) &= \sum_{i=1}^n s(C_{E_i}(\omega), f_i) \quad [\text{Definition 3}] \\ &= \sum_{i=1}^n -\varphi(f_i) - \varphi'(f_i)(C_{E_i}(\omega) - f_i) \quad [\text{Proposition 2}] \\ &= d_\Phi(\mathbf{v}_\omega, \mathbf{f}) - \sum_{i=1}^n \varphi(C_{E_i}(\omega)) \quad [\text{Definition 6}] \\ &= d_\Phi(\mathbf{v}_\omega, \mathbf{f}) + \sum_{i=1}^n s(C_{E_i}(\omega), C_{E_i}(\omega)) \quad [\text{Equation 3}]. \end{aligned} \quad (5)$$

Now assume that \mathbf{f} is incoherent which, by Proposition 1, means that $\mathbf{f} \notin \text{conv}(V)$. According to Eq. (4) of Proposition 3, there exists a $\mathbf{g} \in \text{conv}(V)$, namely the projection

of \mathbf{f} onto $\text{conv}(\mathbf{V})$, such that $d_\Phi(\mathbf{y}, \mathbf{g}) \leq d_\Phi(\mathbf{y}, \mathbf{f}) - d_\Phi(\mathbf{g}, \mathbf{f})$ for all $\mathbf{y} \in \text{conv}(\mathbf{V})$ and hence, in particular, for $\mathbf{y} \in \mathbf{V}$. Since $d_\Phi(\mathbf{g}, \mathbf{f}) > 0$ this proves part (b) of Theorem 1.

To prove part (a) first note that weak dominance of \mathbf{f} by \mathbf{g} means that $d_\Phi(\mathbf{v}_\omega, \mathbf{g}) \leq d_\Phi(\mathbf{v}_\omega, \mathbf{f})$ for all $\mathbf{v}_\omega \in \mathbf{V}$, by Eq. (5). In this case, $d_\Phi(\mathbf{y}, \mathbf{g}) \leq d_\Phi(\mathbf{y}, \mathbf{f})$ for all $\mathbf{y} \in \text{conv}(\mathbf{V})$, since $d_\Phi(\mathbf{y}, \mathbf{g}) - d_\Phi(\mathbf{y}, \mathbf{f})$ depends linearly on \mathbf{y} . If \mathbf{f} is coherent, $\mathbf{f} \in \text{conv}(\mathbf{V})$ by Proposition 1, and hence $d_\Phi(\mathbf{f}, \mathbf{g}) \leq d_\Phi(\mathbf{f}, \mathbf{f}) = 0$. This implies that $\mathbf{g} = \mathbf{f}$.

Unbounded Case.

Next, consider the case when s is unbounded. In this case, the derivative of the corresponding φ from Proposition 2 diverges either at 0 or 1, or at both values, and hence we can not directly apply Proposition 3. Eq. (5) is still valid, with both sides of the equation possibly being $+\infty$. However, if \mathbf{f} lies either in the interior of $[0, 1]^n$, or on a point on the boundary where the derivative of $\Phi(\mathbf{x}) = \sum_i \varphi(x_i)$ does not diverge, an examination of the proof of Proposition 3 shows that the result still applies, as we show now.

If $\nabla\Phi(\mathbf{f})$ is finite, the minimum of $\Phi(\mathbf{y}) - \nabla\Phi(\mathbf{f}) \cdot \mathbf{y}$ over $\mathbf{y} \in \text{conv}(\mathbf{V})$ is uniquely attained at some $\mathbf{g} \in \text{conv}(\mathbf{V})$. Moreover, $\nabla\Phi(\mathbf{g})$ is necessarily finite. Repeating the argument in the proof of Proposition 3 shows that $d_\Phi(\mathbf{y}, \mathbf{g}) \leq d_\Phi(\mathbf{y}, \mathbf{f}) - d_\Phi(\mathbf{g}, \mathbf{f})$ for any $\mathbf{y} \in \text{conv}(\mathbf{V})$, which is the desired inequality needed in the proof of Theorem 1(b). We are thus left with the case in which \mathbf{f} lies on an $(n-1)$ dimensional face of $[0, 1]^n$ where the normal derivative diverges. Consider first the case $n = 1$. Then either $\mathbf{V} = \{0, 1\}$, in which case \mathbf{f} is coherent, or $\mathbf{V} = \{0\}$ or $\{1\}$, in which case it is clear that the unique coherent vector $\mathbf{g} \in \mathbf{V}$ strongly dominates \mathbf{f} .

We now proceed by induction on the dimension n of the forecast \mathbf{f} . In the $(n-1)$ dimensional hypercube, either \mathbf{f} lies inside or on a point of the boundary where the normal derivative of Φ is finite, in which case we have just argued that there exists a $\tilde{\mathbf{g}}$ that is coherent and satisfies $P_s(\omega, \tilde{\mathbf{g}}) < P_s(\omega, \mathbf{f})$ for all ω such that \mathbf{v}_ω lies in the $(n-1)$ dimensional face. In the other case, the induction hypothesis implies that we can find such a $\tilde{\mathbf{g}}$. Note that for all the other ω , $P_s(\omega, \tilde{\mathbf{g}}) = P_s(\omega, \mathbf{f}) = \infty$. Now simply pick an $0 < \epsilon < 1$ and choose $\mathbf{g}_\epsilon = (1 - \epsilon)\tilde{\mathbf{g}} + \epsilon l^{-1} \sum_{i=1}^l \mathbf{v}_i$, where the \mathbf{v}_i denote all the l elements of \mathbf{V} outside the $(n-1)$ -dimensional hypercube. Then $P_s(\omega, \mathbf{g}_\epsilon) < \infty$ for all ω and also, using Lemma 1, $\lim_{\epsilon \rightarrow 0} P_s(\omega, \mathbf{g}_\epsilon) = P_s(\omega, \tilde{\mathbf{g}})$. Hence we can choose ϵ small enough to conclude that $P_s(\omega, \mathbf{g}_\epsilon) < P_s(\omega, \mathbf{f})$ for all $\omega \in \Omega$. This finishes the proof of part (b) in the general case of unbounded s .

To prove part (a) in the general case, we note that if $\mathbf{f} = \sum_i a_i \mathbf{v}_i$ for $\mathbf{v}_i \in \mathbf{V}$ and $a_i > 0$, then necessarily $d_\Phi(\mathbf{v}_i, \mathbf{f}) < \infty$. That is, any coherent \mathbf{f} is a convex combination of $\mathbf{v}_i \in \mathbf{V}$ such that $d_\Phi(\mathbf{v}_i, \mathbf{f}) < \infty$. This follows from the fact that a component of \mathbf{f}

can be 0 only if this component is 0 for all the \mathbf{v}_i 's. The same is true for the value 1. But the $d_\Phi(\mathbf{v}, \mathbf{f})$ can be infinite only if some component of \mathbf{f} is 0 and the corresponding one for \mathbf{v} is 1, or vice versa.

Since $d_\Phi(\mathbf{v}_i, \mathbf{f}) < \infty$ for the \mathbf{v}_i in question, also $d_\Phi(\mathbf{v}_i, \mathbf{g}) < \infty$ by Eq. (5) and the assumption that \mathbf{f} is weakly dominated by \mathbf{g} . Moreover, $d_\Phi(\mathbf{v}_i, \mathbf{g}) - d_\Phi(\mathbf{v}_i, \mathbf{f}) \leq 0$. But $\sum_i \alpha_i (d_\Phi(\mathbf{v}_i, \mathbf{g}) - d_\Phi(\mathbf{v}_i, \mathbf{f})) = d_\Phi(\mathbf{f}, \mathbf{g}) \geq 0$, hence $\mathbf{f} = \mathbf{g}$. \square

6 Proofs of Propositions 1–3

Proof of Proposition 1. Recall that n is the dimension of \mathcal{E} , and that k is the number of elements in V . Let X be the collection of all nonempty sets of form $\bigcap_{i=1}^n E_i^*$, where E_i^* is either E_i or its complement. (X corresponds to the minimal non-empty regions appearing in the Venn diagram of \mathcal{E} .) It is easy to see that:

(a) X partitions Ω .

It is also clear that there is a one-to-one correspondence between X and V with the property that $e \in X$ is mapped to $\mathbf{v} \in V$ such that for all $i \leq n$, $e \subseteq E_i$ iff $v_i = 1$. (Here, v_i denotes the i th component of \mathbf{v} .) Thus, there are k elements in X . We enumerate them as e_1, \dots, e_k , and the corresponding \mathbf{v} by $\mathbf{v}(e_j)$. Plainly, for all $i \leq n$, E_i is the disjoint union of $\{e_j : j \leq k \wedge v(e_j)_i = 1\}$, and hence:

(b) For any measure μ , $\mu(E_i) = \sum_{j=1}^k \mu(e_j) v(e_j)_i$ for all $1 \leq i \leq n$.

For the left-to-right direction of the proposition, suppose that forecast \mathbf{f} is coherent via probability measure μ . Then $f_i = \mu(E_i)$ for all $i \leq n$ and hence by (b), $f_i = \sum_{j=1}^k \mu(e_j) v(e_j)_i$. But the $\mu(e_j)$ are non-negative and sum to one by (a), which shows that $\mathbf{f} \in \text{conv}(V)$.

For the converse, suppose that $\mathbf{f} \in \text{conv}(V)$, which means that there are non-negative α_j 's, with $\sum_j \alpha_j = 1$, such that $\mathbf{f} = \sum_{j=1}^k \alpha_j \mathbf{v}(e_j)$. Let μ be some probability measure such that $\mu(e_j) = \alpha_j$ for all $j \leq k$. By (a) and the assumption about the α_i , it is clear that such a measure μ exists. For all $i \leq n$, $f_i = \sum_{j=1}^k \alpha_j v(e_j)_i = \sum_{j=1}^k \mu(e_j) v(e_j)_i = \mu(E_i)$ by (b), thereby exhibiting \mathbf{f} as coherent. \square

Before giving the proof of Proposition 2, we state and prove the following technical Lemma.

Lemma 1. Let $\varphi : [0, 1] \rightarrow \mathfrak{R}$ be bounded, convex and differentiable on $(0, 1)$. Then the limits $\lim_{p \rightarrow 0, 1} \varphi(p)$ and $\lim_{p \rightarrow 0, 1} \varphi'(p)$ exist, the latter possibly being equal to $-\infty$ at $x = 0$ or $+\infty$ at $x = 1$. Moreover,

$$\lim_{p \rightarrow 0} p\varphi'(p) = \lim_{p \rightarrow 1} \varphi'(p)(1 - p) = 0. \quad (6)$$

Proof of Lemma 1. Since φ is convex, the limits $\lim_{p \rightarrow 0, 1} \varphi(p)$ exist, and they are finite since φ is bounded. Moreover, φ' is a monotone increasing function, and hence also $\lim_{p \rightarrow 0, 1} \varphi'(p)$ exists (but possibly equals $-\infty$ at $x = 0$ or $+\infty$ at $x = 1$). Finally, Eq. (6) follows again from monotonicity of φ' and boundedness of φ , using that $0 = \lim_{p \rightarrow 0} \int_0^p \varphi'(q) dq \leq \lim_{p \rightarrow 0} p\varphi'(p)$, and likewise at $p = 1$. \square

Proof of Proposition 2. Let s be a proper scoring rule. For $0 < p < 1$, let

$$\varphi(p) = -\min_x \{ps(1, x) + (1 - p)s(0, x)\}. \quad (7)$$

By Definition 2(a), the minimum in (7) is achieved at $x = p$, hence $\varphi(p) = -ps(1, p) - (1 - p)s(0, p)$.

As a minimum over linear functions, $-\varphi$ is concave; hence φ is convex. Clearly, φ is bounded (because $s \geq 0$ implies, from (7), that $\varphi \leq 0$, but a convex function can become unbounded only by going to $+\infty$).

The fact that the minimum is achieved uniquely (Def. 2) implies that φ is strictly convex for the following reason. We take $x, y \in (0, 1)$ and $0 < \alpha < 1$ and set $z = \alpha x + (1 - \alpha)y$. Then $\varphi(y) = -y s(1, y) - (1 - y) s(0, y) > -y s(1, x) - (1 - y) s(0, z)$ by uniqueness of the minimizer at $y \neq z$. Similarly, $\varphi(x) = -x s(1, x) - (1 - x) s(0, x) > -x s(1, z) - (1 - x) s(0, z)$. By adding α times the first inequality to $1 - \alpha$ times the second we obtain $\alpha\varphi(y) + (1 - \alpha)\varphi(x) > -z s(1, z) - (1 - z) s(0, z) = \varphi(z)$, which is precisely the statement of strict convexity.

Let $\psi(p) = s(0, p) - s(1, p)$. If φ is differentiable and $\varphi'(p) = \psi(p)$ for all $0 < p < 1$, then (2) is satisfied, as simple algebra shows.

We shall now show that φ is, in fact, differentiable and $\varphi' = \psi$. For any $p \in (0, 1)$ and small enough ϵ , we have

$$\begin{aligned} \frac{1}{\epsilon} (\varphi(p + \epsilon) - \varphi(p)) &= \psi(p) \\ &\quad - \frac{1}{\epsilon} [(p + \epsilon) (s(1, p + \epsilon) - s(1, p)) + (1 - p - \epsilon) (s(0, p + \epsilon) - s(0, p))] . \end{aligned}$$

Since $(p + \epsilon)s(1, x) + (1 - p - \epsilon)s(0, x)$ is minimized at $x = p + \epsilon$ by Definition 2(a), the last term in square brackets is negative. Hence

$$\lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} (\varphi(p + \epsilon) - \varphi(p)) \geq \psi(p),$$

and similarly one shows

$$\lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} (\varphi(p) - \varphi(p - \epsilon)) \leq \psi(p).$$

Since ψ is continuous by Definition 2(b), this shows that φ is differentiable, and hence $\psi = \varphi'$. This proves Eq. (2). Continuity of φ up to the boundary of $[0, 1]$ follows from continuity of s and Lemma 1.

To prove the converse, first note that if φ is bounded and convex on $(0, 1)$, it can be extended to a continuous function on $[0, 1]$, as shown in Lemma 1. Because of strict convexity of φ we have, for $p \in [0, 1]$ and $0 < x < 1$,

$$ps(1, x) + (1 - p)s(0, x) = -\varphi(x) - \varphi'(x)(p - x) \geq -\varphi(p), \quad (8)$$

with equality if and only if $x = p$.

It remains to show that the same is true for $x \in \{0, 1\}$. Consider first the case $x = 0$. We have to show that $ps(1, 0) + (1 - p)s(0, 0) > -\varphi(p)$ for $p > 0$. By continuity of s , Eq. (2) and Lemma 1, we have $s(1, 0) = -\varphi(0) - \lim_{p \rightarrow 0} \varphi'(p)$, while $s(0, 0) = -\varphi(0)$. If $\lim_{p \rightarrow 0} \varphi'(p) = -\infty$, the result is immediate. If $\varphi'(0) := \lim_{p \rightarrow 1} \varphi'(p)$ is finite, we have $-\varphi(0) - p\varphi'(0) > -\varphi(p)$ again by strict convexity of φ .

Likewise, one shows that $ps(1, 1) + (1 - p)s(0, 1) > -\varphi(p)$ for $p < 1$. This finishes the proof that s is a proper scoring rule. \square

Proof of Proposition 3. For fixed $x \in C$, the function $y \mapsto d_\Phi(y, x)$ is strictly convex, and hence achieves a unique minimum at a point π_x in the convex, closed set Z .

Let $y \in Z$. For $0 \leq \epsilon \leq 1$, $(1 - \epsilon)\pi_x + \epsilon y \in Z$, and hence $d_\Phi((1 - \epsilon)\pi_x + \epsilon y, x) - d_\Phi(\pi_x, x) \geq 0$ by the definition of π_x . Since d_Φ is differentiable in the first argument, we can divide by ϵ and let $\epsilon \rightarrow 0$ to obtain

$$0 \leq \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} (d_\Phi((1 - \epsilon)\pi_x + \epsilon y, x) - d_\Phi(\pi_x, x)) = (\nabla \Phi(\pi_x) - \nabla \Phi(x)) \cdot (y - \pi_x).$$

The fact that

$$d_\Phi(y, x) - d_\Phi(\pi_x, x) - d_\Phi(y, \pi_x) = (\nabla \Phi(\pi_x) - \nabla \Phi(x)) \cdot (y - \pi_x)$$

proves the claim. \square

7 Generalizations

7.1 Penalty functions

Theorem 1 holds for a larger class of penalty functions. In fact, one can use different proper scoring rules for every event, and replace (1) by

$$P_s(\omega, \mathbf{f}) = \sum_{i \leq n} s_i(C_{E_i}(\omega), \mathbf{f}_i),$$

where the s_i are possibly distinct proper scoring rules. In this way, forecasts for some events can be penalized differently than others. The relevant Bregman divergence in this case is given by $\Phi(\mathbf{x}) = \sum_i \varphi_i(x_i)$, where φ_i is determined by s_i via (2). Proof of this generalization closely follows the argument given above, so it is omitted. Additionally, by considering more general convex functions Φ our argument generalizes to certain non-additive penalties.

7.2 Generalized scoring rules

7.2.1 Non-uniqueness

If one relaxes the condition of *unique* minimization in Definition 2(a), a weaker form of Theorem 1 still holds. Namely, for any incoherent forecast \mathbf{f} there exists a coherent forecast \mathbf{g} that weakly dominates \mathbf{f} . Strong dominance will not hold in general, as the example of $s(i, x) \equiv 0$ shows.

Proposition 2 also holds in this generalized case, but the function φ need not be *strictly* convex. Likewise, Proposition 3 can be generalized to merely convex (not necessarily strictly convex) Φ but in this case the projection $\pi_{\mathbf{x}}$ need not be unique. Eq. (4) remains valid.

7.2.2 Discontinuity

A generalization that is more interesting mathematically is to discontinuous scoring rules. Proposition 2 can be generalized to scoring rules that satisfy neither the continuity condition in Definition 2 nor unique minimization. (This is also shown in Gneiting and Raftery, 2007).

Proposition 4. Let $s : \{0, 1\} \times [0, 1] \rightarrow [0, \infty]$ satisfy

$$ps(1, x) + (1 - p)s(0, x) \geq ps(1, p) + (1 - p)s(0, p) \quad \forall x, p \in [0, 1]. \quad (9)$$

Then the function $\varphi : [0, 1] \mapsto \Re$ defined by $\varphi(x) = -xs(1, x) - (1-x)s(0, x)$ is bounded and convex. Moreover, there exists a monotone non-decreasing function $\psi : [0, 1] \mapsto \Re \cup \{\pm\infty\}$, with the property that

$$\psi(x) \geq \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} (\varphi(x) - \varphi(x - \epsilon)) \quad \forall x \in (0, 1], \quad (10)$$

$$\psi(x) \leq \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} (\varphi(x + \epsilon) - \varphi(x)) \quad \forall x \in [0, 1), \quad (11)$$

such that

$$s(i, x) = -\varphi(x) - \psi(x)(i - x) \quad \forall x \in (0, 1). \quad (12)$$

Function φ is strictly convex if and only if the inequality (9) is strict for $x \neq p$.

Conversely, if s is of the form (12), with φ bounded and convex and ψ satisfying (10)–(11), then s satisfies (9).

It is a fact (Hardy et al., 1934) that every convex function φ on $[0, 1]$ is continuous on $(0, 1)$ and has a right and left derivative, ψ_R and ψ_L (defined by the right sides of (11) and (10), respectively) at every point (except the endpoints, where it has only a right or left derivative, respectively). Both ψ_R and ψ_L are non-decreasing functions, and $\psi_L(x) \leq \psi_R(x)$ for all $x \in (0, 1)$. Except for countably many points, $\psi_L(x) = \psi_R(x)$, i.e., φ is differentiable. Eqs. (10)–(11) say that $\psi_L(x) \leq \psi(x) \leq \psi_R(x)$.

Note that although $s(0, x)$ and $s(1, x)$ may be discontinuous, the combination $\varphi(x) = -xs(1, x) - (1-x)s(0, x)$ is continuous. Hence, if $s(0, x)$ jumps up at a point x , $s(1, x)$ has to jump down by an amount proportional to $(1-x)/x$.

The proof of Proposition 4 is virtually the same as the proof of Proposition 2, so we omit it.

7.3 Open question

Whether Theorem 1 holds for this generalized notion of a discontinuous scoring rule remains open. The proof of Theorem 1 given here does not extend to the discontinuous case, since for inequality (4) to hold, differentiability of Φ is necessary, in general.

References

- A. Banerjee, X. Guo, and H. Wang. On the optimality of conditional expectation as a Bregman predictor. *IEEE Transactions on Information Theory*, 51(7):2664–2669, 2005.

- J. M. Bernardo and A. F. M. Smith. *Bayesian Theory*. John Wiley & Sons, West Sussex, England, 1994.
- L. M. Bregman. The relaxation method of finding a common point of convex sets and its application to the solution of problems in convex programming. *U. S. S. R. Computational Mathematics and Mathematical Physics*, 78(384):200–217, 1967.
- G. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78:1–3, 1950.
- Y. Censor and S. A. Zenios. *Parallel Optimization: Theory, Algorithms, and Applications*. Oxford University Press, 1997.
- B. de Finetti. *Theory of Probability*, volume 1. John Wiley and Sons, New York, NY, 1974.
- T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, March 2007.
- I. J. Good. Rational decisions. *Journal of the Royal Statistical Society*, 14:107–114, 1952.
- G. H. Hardy, J. E. Littlewood, and G. Pólya. *Inequalities*. Cambridge University Press, 1934.
- J. M. Joyce. A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65: 575–603, 1998.
- D. V. Lindley. Scoring rules and the inevitability of probability. *International Statistical Review*, 50:1–26, 1982.
- L. J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- R. Selten. Axiomatic characterization of the quadratic scoring rule. *Experimental Economics*, 1:43–62, 1998.
- B. Skyrms. *Choice & Chance: An Introduction to Inductive Logic*. Wadsworth, Belmont CA, 2000.

Joel Predd
Rand Corporation
4570 Fifth Avenue, Suite 600
Pittsburgh, PA 15213
jpredd@rand.org

Robert Seiringer
Dept. of Physics
Princeton University
Princeton NJ 08540
rseiring@princeton.edu

Elliott Lieb
Depts. of Mathematics and Physics
Princeton University
Princeton NJ 08540
lieb@princeton.edu

Daniel Osherson
Dept. of Psychology
Princeton University
Princeton NJ 08540
osherson@princeton.edu

Vincent Poor
Dept. of Electrical Engineering
Princeton University
Princeton NJ 08540
poor@princeton.edu

Sanjeev Kulkarni
Dept. of Electrical Engineering
Princeton University
Princeton NJ 08540
kulkarni@princeton.edu